

Network Design and Performance of the System Integration Test, Linked Simulators Phase.

By Mr. Gregory Grundhoffer, Mr. David Brown, and TSgt. Charles Ashton

INTRODUCTION

The System Integration Test (SIT) evaluates the utility of using advance distributed simulation (ADS) to support cost effective testing of an integrated missile weapon/launch aircraft system. The Linked Simulator Phase of the test applies ADS and connects three simulators: the shooter, the target, and the missile. These simulators were located in different installations across the country. This paper discusses the network design and final network architecture. Network personnel tested performance of the network and also came up against unexpected situations. These challenges and their solutions are examined in the paper.

The goals of the Joint Advanced Distributed (JADS) Joint Test Force (JTF) and the Naval Air Warfare Center Weapons Division (NAWCWPNS) networking teams were to create a network able to satisfy the test's Linked Simulator Phase (LSP) requirements and characterize the network performance. This paper discusses some of the networking issues encountered in meeting those goals.

BACKGROUND

The SIT LSP evaluates the ability of ADS to complement and enhance the existing techniques for testing powered and guided weapons delivered against a maneuvering target. The evaluation quantifies the added value of ADS relative to current testing techniques. The missions simulate a shooter aircraft launching an air-to-air missile against a target aircraft. The shooter, target, and missile are represented by geographically separated simulators. The shooter is represented by the F/A-18 Weapon System Support Facility (WSSF) at China Lake, CA. The missile is the AIM-9 Sidewinder Simulation Laboratory (SIMLAB), also at China Lake, CA. The target is represented by the F-14 Weapons System Integration Center (WSIC) at Point Mugu, CA. Test control of this distributed test will be done from the Test Control and Analysis Center (TCAC) located at the JADS JTF in Albuquerque, NM.

THE SIT LSP NETWORK

The network architecture for the SIT LSP is shown in Figure 1. We added a T1 circuit to the existing NAWCWPNS Realtime Network (NRNet) to link the Test Control and Analysis Center at the JADS JTF to the LSP network. A T1 link from Albuquerque to Point Mugu was used for cost reasons rather than bandwidth requirements. We determined that a T1 circuit cost less than a fractional T1 (e.g., a 512 Kbps circuit) for this connection. Two versions of Wellfleet routers were used along with Cisco routers and IDNX routers (running Cisco software) to connect the T1 circuits making up the LSP network.

19990129 071




Figure 1. The SIT LSP Network

DISCUSSION

There are several networking issues that we dealt with during the SIT LSP that may be peculiar to the test and evaluation using ADS. The most important is the network performance monitoring. Other issues are: the network architecture for test and evaluation, data collection considerations, and time synchronization.

NETWORK PERFORMANCE MONITORING

The JADS JTF will collect network performance data to understand the impact the network has on the test itself. If the network has an adverse impact on data quality, the network performance data will be required to make a determination of the effect on the test. The JADS JTF will collect data on several measures of network performance. These include latency, bandwidth utilization, and the number of dropped packets.

Latency

Latency in a network can be caused by the circuit path; routers, hubs and other network equipment; and processes within computers and simulation systems themselves. Latency is hard to measure and there is some disagreement as to what the measurement means. When measuring

network performance, it is sometimes hard to agree upon what comprises the network. In this paper, we will use the broader sense and include the computer systems attached at each node. The time for data to move from a source simulator to the destination simulation can be called the end-to-end latency. Measuring this latency is desirable but difficult.

The JTF is using two methods to examine latency; these are pings and DIS PDU time stamps. Figure 2 depicts the measurement of latency using these two methods. Ping times represent the time of packets sent from one computer to another and the second computer's response. Pings were done prior to the actual test to characterize the network performance. In order to approximate the network burden as it is seen during a test, we use a rate of 50 pings per second with a 144-byte packet. This roughly corresponds to the PDU rates observed on the network and the size of entity state PDUs. Ping times represent round trip times and are used to get a feel for the network performance and the type of latencies to be expected during a test. These ping times represent a simple baseline of the network hardware and links, not including the test simulators. The simulators reside on other subnets that are not accessible because the network interface units (NIUs) do not pass the pings through to the simulators. Even if the pings could reach the simulators, the ping times would not reflect the processing times of the NIUs or the simulators themselves.

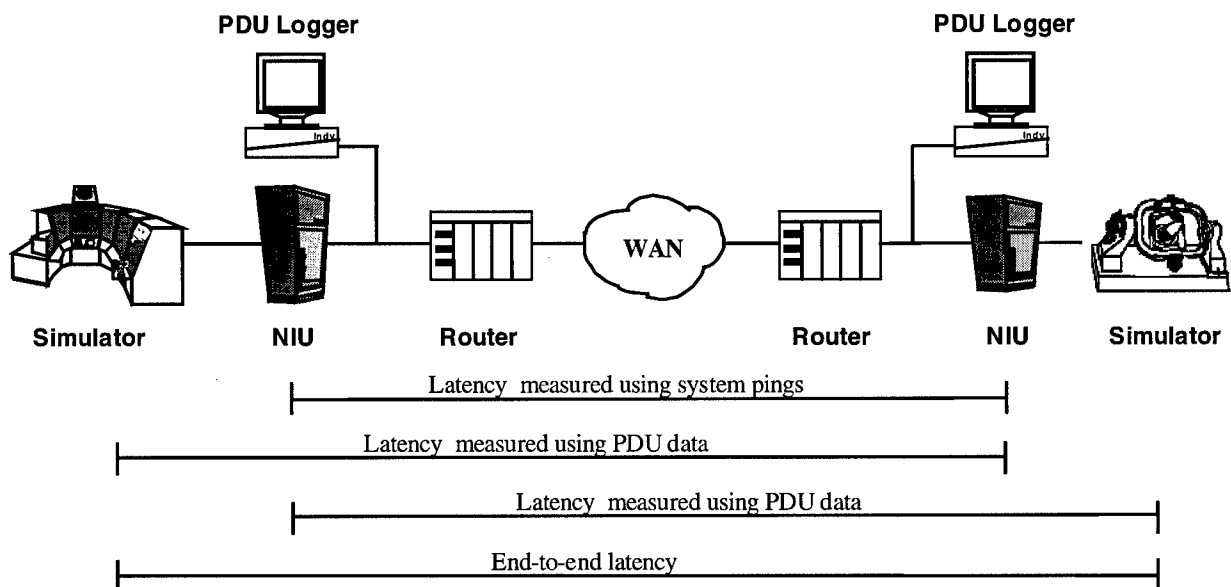


Figure 2. Latency Measurements

The table below shows the preliminary ping data collected prior to the first SIT LSP test in October. At the time this report was written, it was still too early for us to determine the effects of latency on the missile tests. As was previously mentioned, the packet size was 144 bytes and the rate was 50 pings per second.

Table 1. Sample Latency Times

Latency Paths		Percent of Packets Lost	Round-Trip Time (msec.)		
			Minimum	Average	Maximum
TCAC	F-14 WSIC	0	38	38	44
TCAC	F/A-18 WSSF	0	50	51	56
TCAC	SIMLAB	0	49	50	52
F-14 WSIC	F/A-18 WSSF	0	19	20	27
F-14 WSIC	SIMLAB	1	19	20	24
F/A-18 WSSF	SIMLAB	0	7	8	15

The simulators are legacy systems that do not use PDUs to transfer data and require the NIUs to translate PDU data to a format that they can use. The data passed from the simulator to the NIU contains the simulation time. This time is used as the PDU time stamp, as opposed to the time the PDU was created by the NIU. This allows us to look at latency from the time the data was created until it was received. PDU loggers, placed in each lab, stamp the time that the PDU is received. The delta between the creation of a PDU at the simulator and the receipt time recorded by the PDU loggers (refer to Figure 2) represents a portion of the latency within the system. It does not represent the end-to-end latency since it does not include the receiving NIU and simulator processing times. There is not a one-to-one correspondence between PDUs received by a NIU and the updates requested of the NIU by the simulator. That is, the simulators will ask for updates from the NIU at a rate different than the rate the NIU receives PDUs from the network. This precludes the examination of latency using the recorded simulation data. The time a source data item is created can not be easily correlated to the time that the destination simulation uses that information.

Graphs of the latency measurements, of one entity for one run, derived from the PDU log files are shown in Figures 3 and 4. Figure 3 is a histogram representing the latency as measured by the delta from the creation time and the PDU logger receipt time. Figure 4 is a time series of the same data. These data represent an extreme of the latencies looked at so far, not the typically observed latency times. The graphs are presented as examples of the ways we will characterize latency data. Most latencies are well within 100 milliseconds. The cause of the latency anomalies for this entity during this mission is not yet determined.




Figure 3. Example Histogram of PDU Latencies




Figure 4. Example PDU Latency vs Time (milliseconds)

There are several sources of variability in the latency measurements. The DIS community has primarily used UNIX systems. UNIX is not a real-time operating system and thus very accurate time stamping, i.e., millisecond accuracy, is difficult to achieve. This does not usually present a problem for training systems since time synchronization needs are in the order of 100's of milliseconds or greater. This has presented problems for JADS because of testing needs for time accuracy in the milliseconds range. JADS uses COTS tools available from the DIS community for data collection and monitoring because of the high cost of real-time systems and the

development of custom software tools. On a UNIX platform, there may be a delay between the time a PDU packet is received at the network interface and the time it is logged, because other processes may be using CPU cycles. The effect of the UNIX operating system is being investigated by the JTF. Time synchronization also affects the ability to accurately measure latency and will be discussed later.

Bandwidth Utilization

Bandwidth utilization on the LAN segments and the WAN is of interest to all people using ADS. It is of special interest to the test and evaluation community because it is a potential source of data dropouts and increased latency. As bandwidth utilization increases, latency may increase. If bandwidth reaches the limit, latency will increase and will result in data dropouts. In the training community, data dropouts and high latencies are a problem only if the human perceptions of the synthetic world are adversely affected. In the test and evaluation community, where data is the product of testing activities, anything that may adversely affect data quality is of concern.

Theoretical bandwidth on local area Ethernet segments is 10 Megabits per second (Mbps), and is 1.472 Mbps on the T1 links. A UNIX network snooping tool, Netvisualyzer NetGraph, measures bandwidth used on the LAN segments. The Simple Network Management Protocol (SNMP) tool, Cabletron Spectrum, is used to examine router interfaces and determine bandwidth utilization on the WAN circuits. This data collection activity adds traffic to the network. The additional traffic appears to be minimal but the impact has not yet been measured.

The F/A-18 WSSF and SIMLAB simulators are isolated from the DIS network by the NIU which has interfaces on both LANs. Since the simulator traffic is on a physically separate LAN segment than the PDU traffic, the local bandwidth is minimized. At the F-14 WSIC, this separation is not possible because of the physical makeup of the simulator and NIU computers. This results in a much higher bandwidth utilization on that LAN segment. The maximum bandwidth utilization at each site is given below in Table 2.

Table 2. Bandwidth Utilization

Site	Maximum Percent of Bandwidth (10 Mbps) used
JADS TCAC, Albuquerque, NM	1.7
F-14 WSIC, Point Mugu, CA	59.0
F/A-18 WSSF, China Lake, CA	1.1
AIM-9 SIMLAB, China Lake, CA	1.0

We have not yet analyzed data on bandwidth utilization for the T1 circuits but bandwidth does not appear to be a limiting factor for the SIT LSP.

Dropped Packets

To look at dropped packets, we are using the SNMP statistics available from the routers. The routers keep track of the number of packets dropped at each interface and we monitor these from our test control center. Our preliminary results show that no packets are being dropped by the

network devices. At this time, we are unable to determine if any of the dataloggers, simulators, or NIUs are dropping packets because the act of measuring these systems is intrusive and may affect latency and throughput.

NETWORK ARCHITECTURE

Test networks are different from most other networks which makes the design of the network more challenging. Most networks are designed to optimize performance and compatible equipment and configurations are specified. Test programs are not permanent and any network is relatively short-lived. Each test program will connect different ranges and facilities, each site having their own hardware. Often, compatibility of the hardware presents some obstacles. In the SIT LSP, we had two versions of the Wellfleet routers and three versions of the Cisco router software. This presented the most significant design challenge. It impacted not only the network architecture and configuration, but also the configuration of the NIUs.

Since JADS was not the only customer using the F-14 WSIC lab and its local area network (LAN), we were not allowed to change the configuration of that LAN segment. This imposed several problems that we had to overcome. Usually, Internet Protocol (IP) broadcast addressing and router bridging are used to connect all of the nodes participating in a DIS exercise. This could not be done for the SIT LSP because bridging requires a single IP subnet. That is, each LAN needs to have the same broadcast address. This was not possible because the F-14 WSIC IP subnet could not be changed and, therefore, the same subnet could not be used by all of the nodes in the network. Cisco routers have some features (IP helper and UDP Forward Protocol) that allowed us to get around the bridging problems. Each node was given a separate subnet and the IP broadcasts at each lab are re-sent by the Cisco routers to the broadcast addresses at each of the labs. A Wellfleet router at the F/A-18 WSSF was replaced by a Cisco router so that we could implement this architecture. The disadvantage of this approach is that the Cisco routers must send a separate IP (with a PDU) packet, addressed to the remote site's local broadcast address, to each of the three sites. Thus, three packets are forwarded from the Cisco routers for every packet containing PDUs. The Wellfleet routers allow only 1 PDU to be sent to each node.

One option, for future ADS T&E, to avoid incompatibilities is to purchase hardware for each site to ensure everything is compatible. This option may not be feasible because facilities have other users competing for time. Major configuration changes are too disruptive. The F/A-18 WSSF and F-14 WSIC labs supporting our test have many users. We had some control of one of the lab's networks but little control over the other. We had to solve the networking problems while minimizing the impact on the labs and their other users.

DATA COLLECTION REQUIREMENTS

One issue that presented us with a few problems was the data collection requirements (simulator system and network) required to support the SIT LSP. The collection of data is a primary requirement of all test activities and, therefore, must be included in all networking design efforts. Data collection requirements result in additional hardware systems on the network. These data collection systems may collect test system data or network data. While data collection requirements were considered from the beginning of the SIT program, early networking design efforts and tests based upon an earlier program did not take this into account. The early IP addressing scheme had each simulator using point-to-point IP addressing, where PDUs were

addressed to each site that was required to receive them. When the data collection workstations were added to the network, they were unable to see the resultant traffic because they required broadcast addressing, not point-to-point addressing. The bandwidth requirements and load on the NIUs output network interface was too great. The previously mentioned implementation of the Cisco routing configuration at each of the three simulator facilities solved this problem.

TIME SYNCHRONIZATION

To get accurate time tags on the data measurements, the clocks on the systems that record data must be synchronized. This is a topic for a paper itself and will only be briefly discussed here. For the test and evaluation community with highly accurate time sources and real-time systems, time synchronization has not been a problem, but timing concerns are new to the DIS community. DIS has been used primarily by the training community and highly accurate time synchronization was not a requirement. Most of the DIS users that the JTF has talked with do not synchronize time at all.

A UNIX workstation in the TCAC is connected to a GPS receiver to set system time to Universal Coordinated Time (UTC). The PDU loggers at each site use the Network Time Protocol to synchronize to the master clock on a workstation in the TCAC. Again, the fact that UNIX is not a real-time operating system works against us. The clock time on the UNIX workstations drifts from the UTC standard over time and this drift varies from workstation to workstation. Nonetheless, data to date indicates time synchronization in the order of a few milliseconds has been successfully achieved.

SUMMARY

The network design for our test had some unique challenges but they largely represent what other test organizations may experience when connecting multiple facilities to form a network for a test program. NAWCWPNS had significant experience linking their simulation facilities to conduct DIS demonstrations. This experience had both positive and negative impacts on the SIT LSP networking. On the positive side, the NAWCWPNS personnel were able to hit the ground running with experience in connecting systems using DIS. They understood protocols, NIUs, and the networking concepts. The negative aspect of the previous experience was that everyone had a false sense of security. "It cannot be too hard; it has been done before!" The additional requirements of the test were more than expected and had to be accounted for in the network design.

Networks used to support test and evaluation require closer monitoring and baselining than many other networks. This requires more attention to the overall design of the network and the equipment required to support the network. Small anomalies cause no problems in most networks but are generally unacceptable when conducting a test. Testers should verify that the network operates as intended, with all systems on-line, before conducting any test activity.

We expect, due to the nature of networks, that problems will occur due to occasional bit errors, occasional unanticipated traffic, etc., but at this time, we can say that the SIT LSP network seems stable and we expect no impacts on the test, other than latency. When we collect more

data and analyze it further we will determine the impacts of the latency for testers. We hope to share those results at a future ITEA conference.

INTERNET DOCUMENT INFORMATION FORM

A . Report Title: Network Design and Performance of the System Integration Test, Linked Simulators Phase

B. DATE Report Downloaded From the Internet 1/28/99

C. Report's Point of Contact: (Name, Organization, Address, Office Symbol, & Ph #): Joint Advanced Distributed Simulation
Joint Test Force
ATTN: Ann Krause (505) 846-1291
11104 Menaul NE
Albuquerque, NM 87112-2454

D. Currently Applicable Classification Level: Unclassified

E. Distribution Statement A: Approved for Public Release

F. The foregoing information was compiled and provided by:
DTIC-OCA, Initials: VM_ **Preparation Date:** 1/28/99__

The foregoing information should exactly correspond to the Title, Report Number, and the Date on the accompanying report document. If there are mismatches, or other questions, contact the above OCA Representative for resolution.